

REAL *DQDS* FOR THE NONSYMMETRIC TRIDIAGONAL EIGENVALUE PROBLEM *

CARLA FERREIRA[†] AND BERESFORD PARLETT[‡]

Abstract. We present a new transform *triple dqds* to help to compute the eigenvalues of a real tridiagonal matrix C using real arithmetic. The algorithm uses the real *dqds* transform to shift by a real number and *tridqds* to shift by a complex conjugate pair. We present what seems to be a new criteria for splitting the current pair L, U . The algorithm rejects any transform which suffers from excessive element growth and then tries a new transform. Our numerical tests show that the algorithm is about 100 times faster than the Ehrlich-Aberth method of D. A. Bini, L. Gemignani and F. Tisseur. Our code is comparable in performance to a complex *dqds* code and is sometimes 3 times faster.

Key words. LR, *dqds*, unsymmetric tridiagonal matrices

AMS subject classifications. 65F15

1. Introduction. The *dqds* algorithm was introduced in 1994 in [7] as a fast and extremely accurate way to compute all the singular values of a bidiagonal matrix B . This algorithm implicitly performs the Cholesky LR iteration on the tridiagonal matrix $B^T B$ and it is used in LAPACK. However the *dqds* algorithm can also be regarded as executing, implicitly, the LR algorithm applied to any tridiagonal matrix with 1's on the superdiagonal. Our interest is in real matrices which may have complex conjugate pairs of eigenvalues. It is natural to try to retain real arithmetic and yet permit complex shifts of origin. Our analogue of the *double shift* QR algorithm of J. G. F. Francis [12] is the *triple step dqds* algorithm. The purpose of this paper is to explain why 3 steps are needed to derive the algorithm, to explain how we reject transforms with unacceptable element growth and to compare performance with some rival methods. Our conclusion is that this procedure is clearly the fastest method available at the present time.

We say nothing about the need for a tridiagonal eigensolver because this issue is admirably covered in Bini, Gemignani and Tisseur [1]. In fact many parts of [1] have been of great help to us. We also acknowledge the preliminary work on this problem by Z. Wu in [23].

We do not follow Householder conventions except that we reserve capital Roman letters for matrices. Section 2 describes other methods, Section 3 presents standard, but needed, material on LR, *dqds*, double shifts and the implicit L theorem. Section 4 develops our *tridqds* algorithm, Section 5 is our error analysis, Section 6 our splitting, deflation and shift strategy, and Section 7 presents our numerical tests using MATLAB. Finally, Section 8 gives our conclusions and also our ideas about why *tridqds* is only one (important) ingredient for a procedure that must also provide condition numbers and eigenvectors.

2. Other methods.

2.1. 2 steps of LR = 1 step of QR. A frequent exercise for students is to show that for a symmetric positive definite tridiagonal matrix 2 steps of the LR (Cholesky) algorithm produces the same matrix as 1 step of the QR algorithm. Less well known is the article by H. Xu [24] which extends this result when the symmetric matrix is not positive definite. The catch here is that the LR

* The first author is supported by FEDER Funds through "Programa Operacional Factores de Competitividade - COMPETE" and by Portuguese Funds through FCT - "Fundação para a Ciência e a Tecnologia", within the Project PEst-CMAT/UI0013/2011.

[†]Mathematics and Applications Department, University of Minho, 4710-057 Braga (caferrei@math.uminho.pt)

[‡]Department of Mathematics and the Computer Science Division of the Electrical Engineering and Computer Science Department, University of California, Berkeley, California 94720 (parlett@math.berkeley.edu)

transform, if it exists, does not preserve symmetry. The remedy is to regard similarities by diagonal matrices as “trivial”, always available, operations. Indeed, diagonal similarities cannot introduce zeros into a matrix. So, when successful, 2 steps of LR are diagonally similar to one step of QR. Even less well known is a short paper by J. Slemons [20] showing that for a tridiagonal matrix, not necessarily symmetric, 2 steps of LR are diagonally equivalent to 1 step of HR, see [2]. Note that when symmetry disappears then QR is out of the running because it does not preserve the tridiagonal property.

The point of listing these results is to emphasize that 2 steps of LR gives twice as many shift opportunities as 1 step of QR or HR. Thus convergence can be more rapid with LR (or *dqds*) than with QR or HR. This is one of the reasons that *dqds* is faster than QR for computing singular values of bidiagonals. This extra speed is an additional bonus to the fundamental advantage that *dqds* delivers high relative accuracy in all the singular values. The one drawback to *dqds*, for bidiagonals, is that the singular values must be computed in monotone increasing order; QR allows the singular values to be found in any order.

In our case, failure is always possible and so there is no constraint on the order in which eigenvalues are found. The feature of having more opportunities to shift leads us to favor *dqds* over QR and HR. See the list of other methods which follows. We take up the methods in historical order and consider only those that preserve tridiagonal form.

2.2. Cullum’s complex QR algorithm. As part of a program that used the Lanczos algorithm to reduce a given matrix to tridiagonal form in [4], Jane Cullum used the fact that an unsymmetric tridiagonal matrix may always be balanced by a diagonal similarity transformation. She then observed that another diagonal similarity with 1 or i produces a symmetric, but complex, tridiagonal matrix to which the (complex) tridiagonal QR algorithm may be applied. The process is not backward stable because the relation

$$\cos^2 \tau + \sin^2 \tau = 1$$

is not constraint on $|\cos \tau|$ and $|\sin \tau|$ when they are not real. Despite the possibility of breakdown the method proved satisfactory in practice. We have not used it in our comparisons because we are persuaded by 2.1 that it is out performed by the complex *dqds* algorithm, described below.

2.3. Liu’s HR algorithm. In [13] Alex Liu found a variation on the HR algorithm of Angelika Bunse-Gerstner that, in exact arithmetic, is guaranteed not to breakdown - but the price is a temporary increase in bandwidth. This procedure has only been implemented in MAPLE and we do not include it in our comparison.

2.4. Complex *dqds*. In his thesis David Day [5] developed a Lanczos-style algorithm to reduce a general matrix to tridiagonal form and, as with Jane Cullum, needed a suitable algorithm to compute its eigenvalues. He knew of the effectiveness of *dqds* in the symmetric positive definite case and realized that *dqds* extends formally to any tridiagonal that admits triangular factorization. Without positivity the attractive property of achieving high relative accuracy disappears but, despite possible element growth, the error analysis for *dqds* persists: if the transform does not breakdown then tiny well chosen changes in the entries of input L, U (giving \tilde{L}, \tilde{U}) and output \hat{L}, \hat{U} (giving \check{L}, \check{U}) produces an exact relation

$$\check{L}\check{U} = \tilde{L}\tilde{U} - \sigma I$$

with the given shift σ . See Section 5.1. The code uses complex arithmetic because of the possible presence of complex conjugate pairs of eigenvalues. We have wrapped David Day’s complex *dqds* code in a more sophisticated wrapper that chooses suitable shifts after rejecting a transform for excessive element growth.

2.5. Ehrlich-Aberth algorithm. This very careful and accurate procedure was presented by Bini, Gemignani and Tisseur in [1]. It finds the zeros of the characteristic polynomial $p(\cdot)$ and exploits the tridiagonal form to evaluate $p'(z)/p(z)$ for any z . The polynomial solver improves a full set of approximate zeros at each step. Initial approximations are found using a divide-and-conquer procedure that delivers the eigenvalues of the top and bottom halves of the matrix T . The quantity $p'(z)/p(z)$ is evaluated indirectly as $[\text{trace}(zI - T)^{-1}]$ using a QR factorization of $zI - T$. Since T is not altered there is no deflation to assist efficiency. Very careful tests exhibit the method's accuracy - but it is very slow compared to both *dqds*-type algorithms.

3. LR and *dqds*. The reader is expected to have had some exposure to the QR and/or LR algorithms so we will be brief.

3.1. LU factorization. Any $n \times n$ matrix A permits unique triangular factorization $A = LD\tilde{U}$ where L is unit lower triangular, D is diagonal, \tilde{U} is unit upper triangular, if and only if the leading principal submatrices of orders $1, \dots, n-1$ are nonsingular.

In this paper we follow common practice and write $U = D\tilde{U}$ so that the “pivots” (entries of D) lie on U 's diagonal. Throughout this paper any matrix L is unit lower triangular and U is upper triangular.

3.2. LR transform with shift. Note that U is “right” triangular and L is “left” triangular and this explains the standard name LR. For any shift σ let

$$A - \sigma I = LU, \quad (3.1)$$

$$\hat{A} = UL + \sigma I. \quad (3.2)$$

Then \hat{A} is the $\text{LR}(\sigma)$ transform of A . Note that

$$\hat{A} = L^{-1}(A - \sigma I)L + \sigma I = L^{-1}AL.$$

We say that the shift is restored (in contrast to *dqds* - see below). The LR algorithm consists of repeated LR transforms with shifts chosen to enhance convergence to upper triangular form. For the theory see [18, 19, 21, 22].

In contrast to the well known QR algorithm, the LR algorithm can breakdown and can suffer from element growth, $\|L\| \gg \|A\|$, $\|U\| \gg \|A\|$. However LR preserves the banded form of A while QR does not (except for the Hessenberg form).

When a matrix A is represented by its entries then the shift operation $A \rightarrow A - \sigma I$ is trivial. When a matrix is given in factored form the shift operation is not trivial.

3.3. The *dqds* algorithm. From now on we focus on tridiagonal matrices in J -form - entries $(i, i+1)$ are all 1, $i = 1, \dots, n-1$. Throughout this paper all J matrices have this form.

If $J - \sigma I$ permits triangular factorization

$$J - \sigma I = LU$$

then L and U must have the following form

$$L = \begin{bmatrix} 1 & & & & \\ l_1 & 1 & & & \\ & \ddots & \ddots & & \\ & & l_{n-2} & 1 & \\ & & & l_{n-1} & 1 \end{bmatrix}, \quad U = \begin{bmatrix} u_1 & 1 & & & \\ & u_2 & 1 & & \\ & & \ddots & \ddots & \\ & & & u_{n-1} & 1 \\ & & & & u_n \end{bmatrix}. \quad (3.3)$$

It is an attractive feature of LR that

$$UL = \hat{J}$$

is also of J -form. Thus the parameters l_i , $i = 1, \dots, n-1$, and u_j , $j = 1, \dots, n$, determine the matrices L and U above and implicitly define two tridiagonal matrices LU and UL .

The qds algorithm is equivalent to the LR algorithm but no tridiagonal matrices are ever formed. The *progressive* transformation is from L, U to \hat{L}, \hat{U} ,

$$\hat{L}\hat{U} = UL - \sigma I. \quad (3.4)$$

Notice that the shift is not restored and so $\hat{U}\hat{L}$ is not similar to UL .

Equating entries in each side of equation (3.4) gives

$$\begin{aligned} \mathbf{qds}(\sigma) : \quad & \hat{u}_1 = u_1 + l_1 - \sigma; \\ & \mathbf{for} \ i = 1, \dots, n-1 \\ & \quad \hat{l}_i = l_i u_{i+1} / \hat{u}_i \\ & \quad \hat{u}_{i+1} = u_{i+1} + l_{i+1} - \sigma - \hat{l}_i \\ & \mathbf{end for.} \end{aligned}$$

The algorithm qds fails when $\hat{u}_i = 0$ for some $i < n$. When $\sigma = 0$ we write simply qd , not qds .

In 1994 a better way was found to implement $qds(\sigma)$ that had been used by Rutishauser as early as 1955. These are called *differential qd* algorithms. See [15] for more history. This form uses an extra variable d but has compensating advantages.

$$\begin{aligned} \mathbf{dqds}(\sigma) : \quad & d_1 = u_1 - \sigma \\ & \mathbf{for} \ i = 1, \dots, n-1 \\ & \quad \hat{u}_i = d_i + l_i \\ & \quad \hat{l}_i = l_i (u_{i+1} / \hat{u}_i) \\ & \quad d_{i+1} = d_i (u_{i+1} / \hat{u}_i) - \sigma \\ & \mathbf{end for} \\ & \hat{u}_n = d_n. \end{aligned}$$

By definition, $dqd = dqds(0)$.

A word on terminology. In Rutishauser's original work $q_i = u_i$, $e_i = l_i$; the q_i 's were certain *quotients* and the e_i 's were called *modified differences*. In fact the qd algorithm led to the LR algorithm, not vice-versa. The reader can find more information concerning $dqds$ in [15, 16]

One virtue of the $dqds$ and QR transforms is that they work on the whole matrix so that large eigenvalues are converging near the top, albeit slowly, while the small ones are being picked off at the bottom.

We summarize some advantages and disadvantages of the factored form.

Advantages of the factored form

1. L, U determines the entries of J to greater than working-precision accuracy because the addition and multiplication of l 's and u 's is implicit. Thus, for instance, the (i, i) entry of J is given by $l_{i-1} + u_i$ implicitly but $fl(l_{i-1} + u_i)$ explicitly.
2. Singularity of J is detectable by inspection when L and U are given, but only by calculation from J . So, LU reveals singularity, J does not.
3. LU defines the eigenvalues better than J does (usually). There is more on this in [6].

4. Solution of $Jx = b$ takes half the time when L and U are available.

Disadvantages of the factored form

The mapping $J, \sigma \mapsto L, U$ is not everywhere defined for all pairs J, σ and can suffer from element growth. This defect is not as serious as it was when the new transforms were written over the old ones. For tridiagonals we can afford to double the storage and map L, U into different arrays \widehat{L}, \widehat{U} . Then we can decide whether or not to accept \widehat{L}, \widehat{U} and only then would L and U be overwritten. So the difficulty of excessive element growth has been changed from disaster to the non-trivial but less intimidating one of, after rejecting a transform, choosing a new shift that will not spoil convergence and will not cause another rejection.

Now we turn to our main question of $dqds(\sigma)$: how can complex shifts be used without having to use complex arithmetic? This question has a beautiful answer for QR and LR iterations.

3.4. Double shift LR algorithm. We use the J, L and U notation from the previous section. Consider two steps of the LR algorithm with shifts σ_1 and σ_2 ,

$$\begin{aligned} J_2 - \sigma_1 I &= L_2 U_2 \\ J_3 &= U_2 L_2 + \sigma_1 I \\ J_3 - \sigma_2 I &= L_3 U_3 \\ J_4 &= U_3 L_3 + \sigma_1 I. \end{aligned}$$

Then

$$J_4 = \mathcal{L}^{-1} J_2 \mathcal{L} \tag{3.5}$$

with

$$\mathcal{L} = L_2 L_3, \quad \mathcal{U} = U_3 U_2$$

and

$$\begin{aligned} \mathcal{L}\mathcal{U} &= L_2(J_3 - \sigma_2 I)U_2 \\ &= L_2(U_2 L_2 + \sigma_1 I)U_2 - \sigma_2 L_2 U_2 \\ &= L_2 U_2 [L_2 U_2 + (\sigma_1 - \sigma_2)I] \\ &= (J_2 - \sigma_1 I)(J_2 - \sigma_2 I) \\ &= J_2^2 - (\sigma_1 + \sigma_2)J_2 + \sigma_1 \sigma_2 I =: M \end{aligned} \tag{3.6}$$

Suppose that J_2 is real and σ_1 is complex. Then J_4 will be real if, and only if, $\sigma_2 = \bar{\sigma}_1$. The reason is that M is real, so that \mathcal{L} and \mathcal{U} are real and, by (3.5), J_4 is the product of real matrices. Note however that L_2, U_2, L_3, U_3 are all complex. Fortunately it is possible to compute J_4 from J_2 without using J_3 . This depends on the following result.

THEOREM 3.1. [IMPLICIT L THEOREM] *If H_1 and H_2 are unreduced upper Hessenberg matrices and $H_2 = L^{-1}H_1L$, where L is unit lower triangular, then H_2 and L are completely determined by H_1 and column 1 of L , Le_1 . We omit the proof.*

The clever application to J_2 and J_4 is to observe that column 1 of M ,

$$Me_1 = \mathcal{L}\mathcal{U}e_1 = \mathcal{L}e_1 u_{11}, \quad u_{11} = m_{11},$$

is proportional to column 1 of \mathbf{L} and has only three nonzero entries below the diagonal because J_2 is tridiagonal. Now choose

$$\mathcal{L}_1 = I + \mathbf{m}\mathbf{e}_1^T$$

where

$$\mathbf{m} = [0 \quad m_{21}/m_{11} \quad m_{31}/m_{11} \quad 0 \quad \dots \quad 0]^T$$

and perform an explicit similarity transform on J_2 ,

$$\mathcal{L}_1^{-1} J_2 \mathcal{L}_1 =: K.$$

Observe that K is not tridiagonal. In the 6×6 case

$$K = \begin{bmatrix} x & 1 & & & & \\ x & x & 1 & & & \\ + & x & x & 1 & & \\ + & & x & x & 1 & \\ & & & x & x & 1 \\ & & & & x & x \end{bmatrix}.$$

Next we apply a sequence of elementary similarity transformations such that each transformation pushes the 2×1 bulge one row down and one column to the right. Finally the bulge is chased off the bottom to restore the J -form. In exact arithmetic, the implicit L theorem ensures that this technique of *bulge chasing* gives

$$J_4 = (\mathcal{L}_1 \dots \mathcal{L}_{n-1})^{-1} J_2 (\mathcal{L}_1 \dots \mathcal{L}_{n-1}) \quad \text{and} \quad \mathcal{L} = \mathcal{L}_1 \dots \mathcal{L}_{n-1}.$$

4. Triple *dqds* algorithm.

4.1. Connection to LR algorithm. In figure 4.1 we examine the double shift LR transform derived in section 3.4 but with a significant difference. Instead of J_2 being an arbitrary real matrix in J -form, we assume that it is given to us in the form $U_1 L_1$ obtained from one step of the LR algorithm with shift 0 from real J_1 .




FIGURE 4.1. Double shift LR and three steps of *dqds*

The crucial observation is that, along the bottom line L_2, U_2, L_3, U_3 are all complex and so it requires 3 *dqds* steps to go from real L_1, U_1 to real L_4, U_4 . Moreover the non-restoring shifts in *dqds* are

$$\sigma - 0, \quad \bar{\sigma} - \sigma, \quad 0 - \bar{\sigma}.$$

Here is another way of seeing the relation between LR and $dqds$:

$$\begin{aligned} \begin{cases} J_1 = & L_1 U_1 \\ J_2 = & U_1 L_1 \end{cases} \\ \\ \begin{cases} J_2 - \sigma I = & L_2 U_2 \\ J_3 = & U_2 L_2 + \sigma I \end{cases} & \quad L_2 U_2 = U_1 L_1 - \sigma I \\ \\ \begin{cases} J_3 - \bar{\sigma} I = & L_3 U_3 \\ J_4 = & U_3 L_3 + \bar{\sigma} I \end{cases} & \quad L_3 U_3 = (U_2 L_2 + \sigma I) - \bar{\sigma} I \\ \\ \begin{cases} J_4 = & L_4 U_4 \\ \dots \end{cases} & \quad L_4 U_4 = (U_3 L_3 + \bar{\sigma} I) - 0I \end{aligned}$$

Recall from the previous section that the double LR algorithm can work with complex shifts in real arithmetic by *bulge chasing*. The rest of this section develops a form of bulge chasing for the $dqds$ algorithm.

4.2. 3 steps of $dqds$. In contrast to a single $dqds$ step our *triple $dqds$* restores the shifts. Recall from (3.5) in section 3.4 that

$$L_4 U_4 = J_4 = \mathcal{L}^{-1} J_2 \mathcal{L} = \mathcal{L}^{-1} U_1 L_1 \mathcal{L} \quad (4.1)$$

and, since $\sigma_1 = \sigma$ and $\sigma_2 = \bar{\sigma}$, matrix M in (3.6) is given by

$$M = (U_1 L_1)^2 - 2(\Re \sigma_1) U_1 L_1 + |\sigma_1|^2 I. \quad (4.2)$$

The idea is to transform U_1 into L_4 and L_1 into U_4 by bulge chasing in each matrix,

$$L_4 U_4 = \underbrace{\mathcal{L}^{-1} U_1}_{\text{bulge}} \underbrace{L_1 \mathcal{L}}_{\text{bulge}}.$$

Notice that we need to transform an upper bidiagonal into a lower bidiagonal and vice-versa. From the uniqueness of the LU factorization, when it exists, it follows that there is a unique hidden matrix X such that

$$L_4 = \mathcal{L}^{-1} U_1 X^{-1}, \quad X L_1 \mathcal{L} = U_4.$$

For more on X see [14]. The matrix \mathcal{L} is given, from section 3.4 as a product

$$\mathcal{L} = \mathcal{L}_1 \dots \mathcal{L}_{n-1} \mathcal{L}_n$$

($\mathcal{L}_n = I$) and we will gradually construct the matrix X in corresponding factored form $X_n, \dots, X_2 X_1$. In fact we will write each X_i as a product

$$X_i = Y_i Z_i.$$

The details are quite complicated.

4.2.1. Chasing the bulges. Starting with the factors L_1, U_1 and the shift σ , we normalize column 1 of M in (4.2) to form \mathcal{L}_1 , spoil the bidiagonal form with

$$\underbrace{\mathcal{L}_1^{-1}U_1}_{\text{}} \underbrace{L_1\mathcal{L}_1}_{\text{}}$$

and at each *minor* step $i, i = 1, \dots, n$, matrices Z_i, \mathcal{L}_i and Y_i are chosen to chase the bulges. After n minor steps, we obtain L_4 and U_4 ,

$$\begin{aligned} L_4U_4 &= \underbrace{\mathcal{L}_n^{-1} \dots \mathcal{L}_1^{-1}U_1Z_1^{-1}Y_1^{-1} \dots Z_n^{-1}Y_n^{-1}}_{\text{}} \underbrace{Y_nZ_n \dots Y_1Z_1L_1\mathcal{L}_1 \dots \mathcal{L}_n}_{\text{}} \\ &= \underbrace{\mathcal{L}_n^{-1} \dots \mathcal{L}_1^{-1}U_1X_1^{-1} \dots X_n^{-1}}_{\text{}} \underbrace{X_n \dots X_1L_1\mathcal{L}_1 \dots \mathcal{L}_n}_{\text{}} \\ &= \underbrace{\mathcal{L}^{-1}U_1X^{-1}}_{\text{}} \underbrace{XL_1\mathcal{L}}_{\text{}} \end{aligned}$$

Conceptually we create two work arrays F and G . Initially,

$$F = U_1, \quad G = L_1$$

and, finally,

$$F = L_4, \quad G = U_4.$$

For a complex shift σ , the triple *dqds* algorithm has the following matrix formulation:

tridqds($\sigma, \bar{\sigma}$) :

$$\begin{aligned} &F = U_1; \quad G = L_1 \\ &F = FZ_1^{-1}; \quad G = Z_1G \\ &F = \mathcal{L}_1^{-1}F; \quad G = G\mathcal{L}_1 \quad \text{[form } \mathcal{L}_1 \text{ using (4.2)]} \\ &F = FY_1^{-1}; \quad G = Y_1G \\ \\ &\textbf{for } i = 2, \dots, n-3 \\ &\quad F = FZ_i^{-1}; \quad G = Z_iG \\ &\quad F = \mathcal{L}_i^{-1}F; \quad G = G\mathcal{L}_i \\ &\quad F = FY_i^{-1}; \quad G = Y_iG \quad \text{[} Z_i \text{ with one, } \mathcal{L}_i \text{ with two and } Y_i \text{ with three} \\ &\textbf{end for} \quad \text{nonzero off-diagonal entries]} \\ \\ &\% \text{ step } n-2 \\ &F = FZ_{n-2}^{-1}; \quad G = Z_{n-2}G \\ &F = \mathcal{L}_{n-2}^{-1}F; \quad G = G\mathcal{L}_{n-2} \\ &F = FY_{n-2}^{-1}; \quad G = Y_{n-2}G \quad \text{[} Y_{n-2} \text{ with two nonzero off-diagonal entries]} \\ \\ &\% \text{ step } n-1 \\ &F = FZ_{n-1}^{-1}; \quad G = Z_{n-1}G \\ &F = \mathcal{L}_{n-1}^{-1}F; \quad G = G\mathcal{L}_{n-1} \\ &F = FY_{n-1}^{-1}; \quad G = Y_{n-1}G \quad \text{[} Y_{n-1} \text{ and } \mathcal{L}_{n-1} \text{ with one nonzero off-diagonal entry]} \\ \\ &\% \text{ step } n \\ &\mathcal{L}_n = I; \quad Y_n = I \\ &F = FZ_n^{-1}; \quad G = Z_nG \quad \text{[} Z_n \text{ diagonal]} \\ &L_4 = F; \quad F_4 = G \end{aligned}$$

4.2.2. Details of $tridqds$. In this section we will go into the details of the $tridqds$ algorithm described in the previous section. Consider L_1 with subdiagonal entries l_1, \dots, l_{n-1} and U_1 with diagonal entries u_1, \dots, u_n , as defined in Section 3.3, and consider matrices L_4 and U_4 with subdiagonal entries $\hat{l}_1, \dots, \hat{l}_{n-1}$ and diagonal entries $\hat{u}_1, \dots, \hat{u}_n$, respectively.

For each iteration of $tridqds$, at the beginning of a minor step i , $i = 2, \dots, n-2$, the active 4×4 windows of F and G are

$$F = \begin{bmatrix} \ddots & & & & & \\ \ddots & 1 & & & & \\ & \hat{l}_{i-1} & u_i & 1 & & \\ & + & & u_{i+1} & 1 & \\ & + & & & u_{i+2} & \ddots \\ & & & & & \ddots \end{bmatrix}, \quad G = \begin{bmatrix} \ddots & & \ddots & & & \\ & \hat{u}_{i-1} & 1 & & & \\ & & * & & & \\ & & + & 1 & & \\ & & + & l_{i+1} & 1 & \\ & & & & \ddots & \ddots \end{bmatrix}. \quad (4.3)$$

Each minor step i , $i = 2, \dots, n-3$, consists of the following 3 parts.

Minor step i

- a) $F \leftarrow FZ_i^{-1}$ puts 0 into $F_{i,i+1}$ and 1 into $F_{i,i}$
 $G \leftarrow Z_i G$ turns $G_{i,i+1}$ into 1

$$Z_i^{-1} = \begin{bmatrix} \ddots & & & & & \\ & 1 & & & & \\ & & \frac{1}{u_i} & -\frac{1}{u_i} & & \\ & & 0 & 1 & & \\ & & & & 1 & \\ & & & & & \ddots \end{bmatrix}, \quad Z_i = \begin{bmatrix} \ddots & & & & & \\ & 1 & & & & \\ & & u_i & 1 & & \\ & & 0 & 1 & & \\ & & & & 1 & \\ & & & & & \ddots \end{bmatrix},$$

$$FZ_i^{-1} = \begin{bmatrix} \ddots & & & & & \\ \ddots & 1 & & & & \\ & \hat{l}_{i-1} & 1 & 0 & & \\ & + & & u_{i+1} & 1 & \\ & + & & & u_{i+2} & \ddots \\ & & & & & \ddots \end{bmatrix}, \quad Z_i G = \begin{bmatrix} \ddots & & \ddots & & & \\ & \hat{u}_{i-1} & 1 & & & \\ & & * & 1 & & \\ & & + & 1 & & \\ & & + & l_{i+1} & 1 & \\ & & & & \ddots & \ddots \end{bmatrix}.$$

- b) $F \leftarrow \mathcal{L}_i^{-1} F$ puts 0 in $F_{i+1,i-1}$ and $F_{i+2,i-1}$
 $G \leftarrow G\mathcal{L}_i$ defines \hat{u}_i and creates 3 nonzeros below it

$$\mathcal{L}_i^{-1} = \begin{bmatrix} \ddots & & & & \\ & \ddots & & & \\ & & \ddots & & \\ & & & 1 & \\ & & & * & 1 \\ & & & * & \\ & & & & 1 & \\ & & & & & \ddots \end{bmatrix} = I + \mathbf{x}\mathbf{e}_i^T, \quad \mathcal{L}_i = I - \mathbf{x}\mathbf{e}_i^T,$$

$$\mathcal{L}_i^{-1}F = \begin{bmatrix} \ddots & & & & \\ & \ddots & & & \\ & & 1 & & \\ & & \hat{l}_{i-1} & 1 & \\ & & & * & u_{i+1} & 1 \\ & & & + & & u_{i+2} & 1 \\ & & & & & & \ddots \end{bmatrix}, \quad G\mathcal{L}_i = \begin{bmatrix} \ddots & \ddots & & & \\ & \hat{u}_i & 1 & & \\ & + & 1 & & \\ & + & l_{i+1} & 1 & \\ & + & & l_{i+2} & 1 \\ & & & & \ddots & \ddots \end{bmatrix}.$$

- c) $G \leftarrow Y_i G$ puts 0 in $G_{i+1,i}$, $G_{i+2,i}$ and $G_{i+3,i}$
 $F \leftarrow FY_i^{-1}$ creates \hat{l}_i and puts 2 nonzeros below it

$$Y_i^{-1} = \begin{bmatrix} \ddots & & & & \\ & \ddots & & & \\ & & 1 & & \\ & & * & 1 & \\ & & * & & 1 \\ & & * & & & 1 \\ & & & & & & \ddots \end{bmatrix} = I + \mathbf{y}\mathbf{e}_i^T, \quad Y_i = I - \mathbf{y}\mathbf{e}_i^T,$$

$$FY_i^{-1} = \begin{bmatrix} \ddots & & & & \\ & \ddots & & & \\ & & 1 & & \\ & & \hat{l}_i & u_{i+1} & 1 \\ & & + & & u_{i+2} & 1 \\ & & + & & & u_{i+3} & \ddots \\ & & & & & & \ddots \end{bmatrix}, \quad Y_i G = \begin{bmatrix} \ddots & \ddots & & & \\ & \hat{u}_i & 1 & & \\ & & * & & \\ & & + & 1 & \\ & & + & l_{i+2} & 1 \\ & & & & \ddots & \ddots \end{bmatrix}.$$

The result of this minor step is that the active windows of F and G shown in (4.3) have been moved down and to the right by one place. See Appendix A for more details on the practical implementation.

Naturally steps $1, n-2, n-1, n$ are slightly different and may be found on pp. 147-157 of [8].

4.3. Operation count for $tridqds$. In this section we will see how three steps of simple $dqds$ algorithm compares with one step of $tridqds$ in what respects to the number of floating point operations required.

Here is the inner loop of $tridqds$. See Appendix B.

```

tridqds( $\sigma, \bar{\sigma}$ ) :
  for  $i = 2, \dots, n - 3$ 
     $x_r = x_r * u_i + y_r$ 
     $x_l = -x_l * (1/\hat{l}_{i-1}); \quad y_l = -y_l * (1/\hat{l}_{i-1});$ 
     $\hat{u}_i = x_r - x_l;$ 
     $x_r = y_r - x_l; \quad y_r = z_r - y_l - x_l * l_{i+1}; \quad z_r = -y_l * l_{i+2}$ 
     $x_r = x_r * (1/\hat{u}_i); \quad y_r = y_r * (1/\hat{u}_i); \quad z_r = z_r * (1/\hat{u}_i)$ 
     $\hat{l}_i = x_l + y_r + x_r * u_{i+1}$ 
     $x_l = y_l + z_r + y_r * u_{i+2}; \quad y_l = z_r * u_{i+3}$ 
     $x_r = 1 - x_r; \quad y_r = l_{i+1} - y_r; \quad z_r = -z_r$ 
  end for

```

A good compiler recognizes common subexpressions.

In contrast,

```

dqds( $\sigma$ ) :  $d_1 = u_1 - \sigma$ 
  for  $i = 1, \dots, n - 1$ 
     $\hat{u}_i = d_i + l_i$ 
     $\hat{l}_i = l_i(u_{i+1}/\hat{u}_i)$ 
     $d_{i+1} = d_i(u_{i+1}/\hat{u}_i) - \sigma$ 
  end for
   $\hat{u}_n = d_n.$ 

```

In practice, each d_{i+1} may be written over its predecessor in a single variable d and, if the common subexpression u_{i+1}/\hat{u}_i is recognized, then only one division is needed if we use an auxiliary variable.

Table 4.1 below shows that the operation count of one step of $tridqds$ is comparable to three steps of $dqds$ (table expresses only the number of floating point operations in the inner loops).

	<i>tridqds</i>	3 <i>dqds</i> steps
Divisions	2	3
Multiplications	11	6
Additions	5	3
Subtractions	6	3
Assignments	16	12
Auxiliary variables	5	2

TABLE 4.1
Operation count of $tridqds$ and 3 $dqds$ steps

But to make three steps of $dqds$ equivalent to $tridqds$ we have to consider $dqds$ in complex arithmetic and the total cost is raised by a factor of about 4. Thus, in complex arithmetic, three steps of $dqds$ are much more expensive than one step of $tridqds$.

5. Error analysis. We turn to the effect of finite precision arithmetic on our algorithms. First consider the dqds algorithm.

5.1. dqds. In the absence of over/underflow the algorithm enjoys the so-called *mixed relative stability* property.

THEOREM 5.1. *Let $\text{dqds}(\sigma)$ map L, U into computed \widehat{L}, \widehat{U} with no division by zero, over/underflow. Then well chosen small relative changes in the entries of both input and output matrices, of at most 3 ulps each, produces new matrices, one pair mapped into the other, in exact arithmetic, by $\text{dqds}(\sigma)$.*

See the diagram in Figure 5.1. The remarkable feature here is that huge element growth does not impair the result. However this useful property does not guarantee that dqds returns accurate eigenvalues. See [7, 15]. For that, an extra requirement is needed such as positivity of all the parameters u_j, l_j in the computation. This is the case for the eigenvalues of $B^T B$ where B is upper bidiagonal.

What can be said in our case? We quote a result that is established by Yao Yang in his dissertation [25] and appears in [15]. The clever idea is not to look at the L and U separately but to study their exact product $J = LU$.




FIGURE 5.1. Effects of roundoff for dqds

THEOREM 5.2. [Y. Yang] *If $\text{dqds}(\sigma)$ maps L, U into \widehat{L}, \widehat{U} (with no division by 0, overflow/underflow) in the standard model of floating point arithmetic then there is a unique pair \check{L}, \check{U} such that, in exact arithmetic, $\text{dqds}(\sigma)$ maps \check{L}, \check{U} into \widehat{L}, \widehat{U} . Moreover, the associated tridiagonals satisfy, element by element,*

$$\begin{aligned}
 |\text{offdiag}(\check{J}) - \text{offdiag}(J)| &< 2\varepsilon |\text{offdiag}(J)| \\
 |\text{diag}(\check{J}) - \text{diag}(J)| &< \varepsilon \left(2|\mathbf{u}| + |\sigma||\mathbf{1}| + |\hat{\mathbf{l}}| + |\hat{\mathbf{u}}| + 2|\mathbf{d}| \right)
 \end{aligned}$$

where ε is the roundoff unit.

This result is Corollary 3 in Section 9 of [15]. It shows that it is only the diagonal of J that suffers large backward error in the case of element growth. Since $\hat{\mathbf{u}} = \mathbf{d} + \hat{\mathbf{l}}$ the last inequality may be written as

$$|\text{diag}(\check{J}) - \text{diag}(J)| < \varepsilon \left(2|\mathbf{u}| + |\mathbf{l}| + |\sigma||\mathbf{1}| + |\hat{\mathbf{l}}| + 3|\mathbf{d}| \right).$$

Recall that $d_i^{-1} = [(UL)^{-1}]_{ii}$, $i = 1, \dots, n$. Thus the indices vulnerable to large backward error belong to any very small entries $[(UL)^{-1}]_{ii}$. For this reason we reject \widehat{L}, \widehat{U} when, element by element,

$$|\sigma||\mathbf{1}| + |\hat{\mathbf{l}}| + 3|\mathbf{d}| > 1000(|\mathbf{u}| + |\mathbf{l}|). \quad (5.1)$$

Recall that the error analysis is worst case. Recall also that the effect of a tiny \hat{u}_k disappears for $i > k + 1$.

5.2. *tridqds*. There are too many intermediate variables in this algorithm to permit a successful mixed error analysis. However each minor step in the algorithm consists of 3 elementary similarity transformations on work matrices F, G or G, F . See ... in Section 4.2.1. Recall that an elementary matrix here is of the form $I + \mathbf{v}e_j^T$, with inverse $I - \mathbf{v}e_j^T$, and \mathbf{v} has at most 3 nonzero entries. So we examine the condition number of these 3 similarity transforms. Consult Appendix A to follow the details.

- The active part of Z_i is

$$\begin{bmatrix} u_i & 1 \\ 0 & 1 \end{bmatrix} \quad \text{and} \quad \text{cond}(Z_i) \simeq \max \{|u_i|, |u_i|^{-1}\}.$$

- The active part of \mathcal{L}_i is

$$\begin{bmatrix} 1 & & \\ -x_l/\hat{l}_{i-1} & 1 & \\ -y_l/\hat{l}_{i-1} & 0 & 1 \end{bmatrix} \quad \text{and} \quad \text{cond}(\mathcal{L}_i) \simeq 1 + \left(\frac{x_l}{\hat{l}_{i-1}}\right)^2 + \left(\frac{y_l}{\hat{l}_{i-1}}\right)^2.$$

- The active part of Y_i is

$$\begin{bmatrix} 1 & & & \\ -x_r/\hat{u}_i & 1 & & \\ -y_r/\hat{u}_i & 0 & 1 & \\ -z_r/\hat{u}_i & 0 & 0 & 1 \end{bmatrix} \quad \text{and} \quad \text{cond}(Y_i) \simeq 1 + \left(\frac{x_r}{\hat{u}_i}\right)^2 + \left(\frac{y_r}{\hat{u}_i}\right)^2 + \left(\frac{z_r}{\hat{u}_i}\right)^2.$$

The variables x_l, y_l, x_r, y_r, z_r are formed from additions and multiplications of previous quantities. Note that u_i is part of the input and so is assumed to be of acceptable size. We see that it is tiny values of \hat{l}_{i-1} and \hat{u}_i that lead to an ill-conditioned similarity at minor step i . In the simple $dqds$ algorithm a small value of \hat{u}_i (relative to u_{i+1}) leads to a large value of \hat{l}_{i+1} and d_{i+1} . In *tridqds* the effect of 3 consecutive transforms is more complicated. The message is the same: reject any transform that has more than modest element growth, as determined by (5.1) in the previous section. This challenge calls for further study.

6. Implementation details.

6.1. Deflation ($n \leftarrow n - 1$). Some of our criteria for deflating come from [16], others are new. Consider both matrices UL and LU and the trailing 2×2 blocks,

$$\begin{bmatrix} l_{n-1} + u_{n-1} & 1 \\ l_{n-1}u_n & u_n \end{bmatrix}, \quad \begin{bmatrix} l_{n-2} + u_{n-1} & 1 \\ l_{n-1}u_{n-1} & l_{n-1} + u_n \end{bmatrix}.$$

Deflation ($n \leftarrow n - 1$) removes l_{n-1} as well as u_n . Looking at entry $(n-1, n-1)$ of UL shows that a necessary condition is that l_{n-1} be negligible compared to u_{n-1} ,

$$|l_{n-1}| < \text{tol} \cdot |u_{n-1}|, \quad (6.1)$$

for a certain tolerance tol close to roundoff unit ε .

The (n, n) entries of UL and LU suggest either $u_n + \text{Acshift}$ or $l_{n-1} + u_n + \text{Acshift}$ as eigenvalues. *Acshift* is the accumulated shift (recall that $dqds$ is a non-restoring transform). To make these consistent we require that

$$|l_{n-1}| < \text{tol} \cdot |u_n + \text{Acshift}|. \quad (6.2)$$

Finally we must consider the change $\delta\lambda$ in the eigenvalue λ caused by setting $l_{n-1} = 0$. We estimate $\delta\lambda$ by starting from UL with $l_{n-1} = 0$ and then allowing l_{n-1} to grow. To this end let J be UL with $l_{n-1} = 0$ and $(u_n, \mathbf{y}^T, \mathbf{x})$ be the eigentriple for J . Clearly $\mathbf{y}^T = \mathbf{e}_n^T$. Now we consider perturbation theory with parameter l_{n-1} . The perturbing matrix δJ , as l_{n-1} grows, is

$$l_{n-1}(\mathbf{e}_{n-1} + \mathbf{e}_n u_n) \mathbf{e}_{n-1}^T.$$

By first order perturbation analysis

$$|\delta\lambda| = \frac{|\mathbf{y}^T \delta J \mathbf{x}|}{\|\mathbf{x}\|_2 \|\mathbf{y}\|_2}$$

and $\|\mathbf{y}\|_2 = 1$ in our case. So,

$$|\delta\lambda| = \frac{l_{n-1} \mathbf{e}_n^T (\mathbf{e}_{n-1} + \mathbf{e}_n u_n) \mathbf{e}_{n-1}^T \mathbf{x}}{\|\mathbf{x}\|_2} = \frac{l_{n-1} u_n |x_{n-1}|}{\|\mathbf{x}\|_2}$$

and we use the crude bound $\frac{|x_{n-1}|}{\|\mathbf{x}\|_2} < 1$. So, we let l_{n-1} grow until the change

$$|\delta\lambda| < |l_{n-1} u_n|$$

in eigenvalue $\lambda = u_n$ is no longer acceptable. Our condition for deflation is then

$$|l_{n-1} u_n| < \text{tol} \cdot |\text{Acshift} + u_n|. \quad (6.3)$$

A similar first order perturbation analysis for LU with $l_{n-1} = 0$ will give our last condition for deflation. For the eigentriple $(u_n, \mathbf{y}^T, \mathbf{x})$ we also have $\mathbf{y}^T = \mathbf{e}_n^T$. The perturbing matrix is now

$$l_{n-1} \mathbf{e}_n (\mathbf{e}_{n-1}^T u_{n-1} + \mathbf{e}_n^T)$$

and

$$|\delta\lambda| = \frac{l_{n-1} \mathbf{e}_n^T \mathbf{e}_n (\mathbf{e}_{n-1}^T u_{n-1} + \mathbf{e}_n^T) \mathbf{x}}{\|\mathbf{x}\|_2} = |l_{n-1}| \frac{|u_{n-1} x_{n-1} + x_n|}{\|\mathbf{x}\|_2} < |l_{n-1}| (|u_{n-1}| + 1).$$

Finally we require

$$|l_{n-1}| (|u_{n-1}| + 1) < \text{tol} \cdot |\text{Acshift} + u_n|. \quad (6.4)$$

6.2. Splitting and deflation ($n \leftarrow n - 2$). Recall that the implicit L theorem was invoked to justify the *tridqds* algorithm. This result fails if any l_k , $k < n - 1$ vanishes. Consequently, checking for negligible values among the l_k is a necessity, not a luxury for increased efficiency. Consider $J = UL$ in block form

$$\left[\begin{array}{c|c} J_1 & \\ \hline \mu & J_2 \end{array} \right]$$

where $\mu = u_{k+1}l_k$, $k < n - 1$. We can replace μ by 0 when

$$\text{spectrum}(J_1) \cup \text{spectrum}(J_2) = \text{spectrum}(J), \quad \text{to working accuracy.}$$

However we are not going to estimate the eigenvalues of J_1 and J_2 . Instead we create a local criterion for splitting at $(k + 1, k)$ as follows. Focus on the principal 4×4 window of J given by

$$\begin{bmatrix} u_{k-1} + l_{k-1} & 1 & & \\ u_k l_{k-1} & u_k + l_k & 1 & \\ \hline & u_{k+1} l_k & u_{k+1} + l_{k+1} & 1 \\ & & u_{k+2} l_{k+1} & u_{k+2} + l_{k+2} \end{bmatrix}.$$

Now J_1 and J_2 are both 2×2 and our local criterion is

$$\det(J_1) \cdot \det(J_2) = \det(J), \quad \text{to working accuracy.} \quad (6.5)$$

Let us see what this yields. Perform block factorization on J and note that the Schur complement of J_1 in J is

$$J'_2 = J_2 - \begin{bmatrix} 0 & \mu \\ 0 & 0 \end{bmatrix} J_1^{-1} \begin{bmatrix} 0 & 0 \\ 1 & 0 \end{bmatrix}$$

with

$$J_1^{-1} = \frac{1}{\det_1} \begin{bmatrix} u_k + l_k & -1 \\ -u_k l_{k-1} & u_{k-1} + l_{k-1} \end{bmatrix}$$

where

$$\det_1 = \det(J_1) = u_{k-1}(u_k + l_k) + l_{k-1}l_k.$$

Thus

$$J'_2 = \begin{bmatrix} u_{k+1}l_k & u_{k+1} + l_{k+1} \\ u_{k+2}l_{k+1} & u_{k+2} + l_{k+2} \end{bmatrix} - \begin{bmatrix} \mu(u_{k-1} + l_{k-1})/\det_1 & 0 \\ 0 & 0 \end{bmatrix}.$$

Since \det is linear by rows

$$\det(J_2) - \det(J'_2) = \mu(u_{k-1} + l_{k-1})(u_{k+2} + l_{k+2})/\det_1.$$

Our criterion reduces to splitting only when

$$\det(J'_2) = \det(J_2), \quad \text{to working accuracy.}$$

Thus we require

$$|l_k u_{k+1}(u_{k+2} + l_{k+2})(u_{k-1} + l_{k-1})/\det_1| < \text{tol} \cdot |\det(J_2)|.$$

Since

$$\det_2 = \det(J_2) = u_{k+1}(u_{k+2} + l_{k+2}) + l_{k+1}l_{k+2},$$

the criterion for splitting J at $(k+1, k)$ is then

$$|l_k u_{k+1}(u_{k+2} + l_{k+2})(u_{k-1} + l_{k-1})| < tol \cdot |det_1 det_2|. \quad (6.6)$$

Finally, to remove l_k we also need l_k to be negligible compared to u_k ,

$$|l_k| < tol \cdot |u_k|. \quad (6.7)$$

Deflation ($n \leftarrow n - 2$)

We use the same criterion for deflation ($n \leftarrow n - 2$), but because $l_{k+2} = l_n = 0$ there is a common factor det_2 on each side of (6.6). Deflate the trailing 2×2 submatrix when

$$|l_{n-2}| < tol \cdot |u_{n-2}| \quad (6.8)$$

and

$$|l_{n-2}(u_{n-3} + l_{n-3})| < tol \cdot |u_{n-3}(u_{n-2} + l_{n-2}) + l_{n-3}l_{n-2}|. \quad (6.9)$$

We omit the role of *Acshift* here because it makes the situation more complicated. We have to recall that *tridqds* uses restoring shifts and *Acshift* is always real. So, for complex shifts, det_2 is not going to zero. In fact

$$|det_2| \geq |\Im(\lambda)|^2$$

where λ is an eigenvalue of J_2 .

When $n = 3$ these criteria simplify a lot. Both reduce to

$$|l_1| < tol \cdot |u_1|.$$

6.3. Shift strategy. Although *tridqds* may be, and has been, used to compute all the eigenvalues, it seems sensible to include real *dqds*(σ) so that when all eigenvalues are real *tridqds* need not be called.

As with LR, the *dqds* algorithm with no shift gradually forces large entries to the top and brings small entries towards the bottom. Before every transform both l_{n-1} and l_{n-2} are inspected. If

$$|l_{n-1}| < \frac{1}{2^4} \quad \text{and} \quad |l_{n-2}| < \frac{1}{2^4}$$

then the code executes *dqds* transform with the *Wilkinson* shift or a *3dqds* transform with *Francis* shifts depending on the sign of the discriminant.

An unexpected reward for having both transforms available is to cope with a rejected transform. Our strategy is simply to use the other transform with the current shift. More precisely, given a complex shift σ , if *tridqds*($\sigma, \bar{\sigma}$) is rejected we try *dqds*(u_n); if for real τ , *dqds*(τ) is rejected, we try *tridqds*($\tau, \bar{\tau}$). So far, this has not failed.

More generally, an increase in the imaginary part of the shift increases diagonal dominance. At the extreme, consider a pair of pure imaginary shifts $i\mu, -i\mu$, μ positive. The *tridqds* wants $UL + \mu^2 I$ to permit triangular factorization. The bigger is μ the better.

A great attraction of IEEE arithmetic standard is that it allows the symbols *inf* and *NaN*. Thus there is no need for time consuming with if statements in the main loop. At the end of the loops we test for rejection or excessive element growth. We record the number of rejections.

7. Numerical Examples. Those who work with well defined problems have the habit of determining the “true” (or most accurate) solutions and comparing computed values with them to give the error. The condition number of every eigenvalue of a real symmetric matrix is 1, but only in the absolute sense. The relative condition number can vary. In our case even the absolute condition numbers can rise to ∞ and little is known about relative errors.

We have discovered [9] that more often than not the eigenvalues of tridiagonal, and reasonable well balanced, matrices are well determined by an LU or ΔLDL^T representation (Δ is a signature matrix). This is good news but much work remains. Our main focus is on the time it takes to get reasonable approximations, recognizing that we do not know how well the data defines the eigenvalues.

We refer to the Ehrlich-Aberth algorithm (see Section 2.5) as BGT and to our code simply as $tridqds$, although we combine $tridqds$ with real $dqds$ as described in Section 6.3.

Since we compare MATLAB versions of all the codes we acknowledge that the elapsed times are accurate to only about 0.02 seconds. However this is good enough to show the ratios between BGT and the $dqds$ codes. The efficiency of complex $dqds$ is harder to determine. Sometimes the same, sometimes $tridqds$ is 3 times faster.

Since the number of iterations needed for convergence on our (modest) test bed has remained about $4n$, we have not tried for a strategy as sophisticated as the one in [16].

Bessel matrix

Bessel matrices, associated with generalized Bessel polynomials, are nonsymmetric tridiagonals matrices defined by $B_n^{(a,b)} = \text{tridiag}(\beta, \alpha, \gamma)$ with

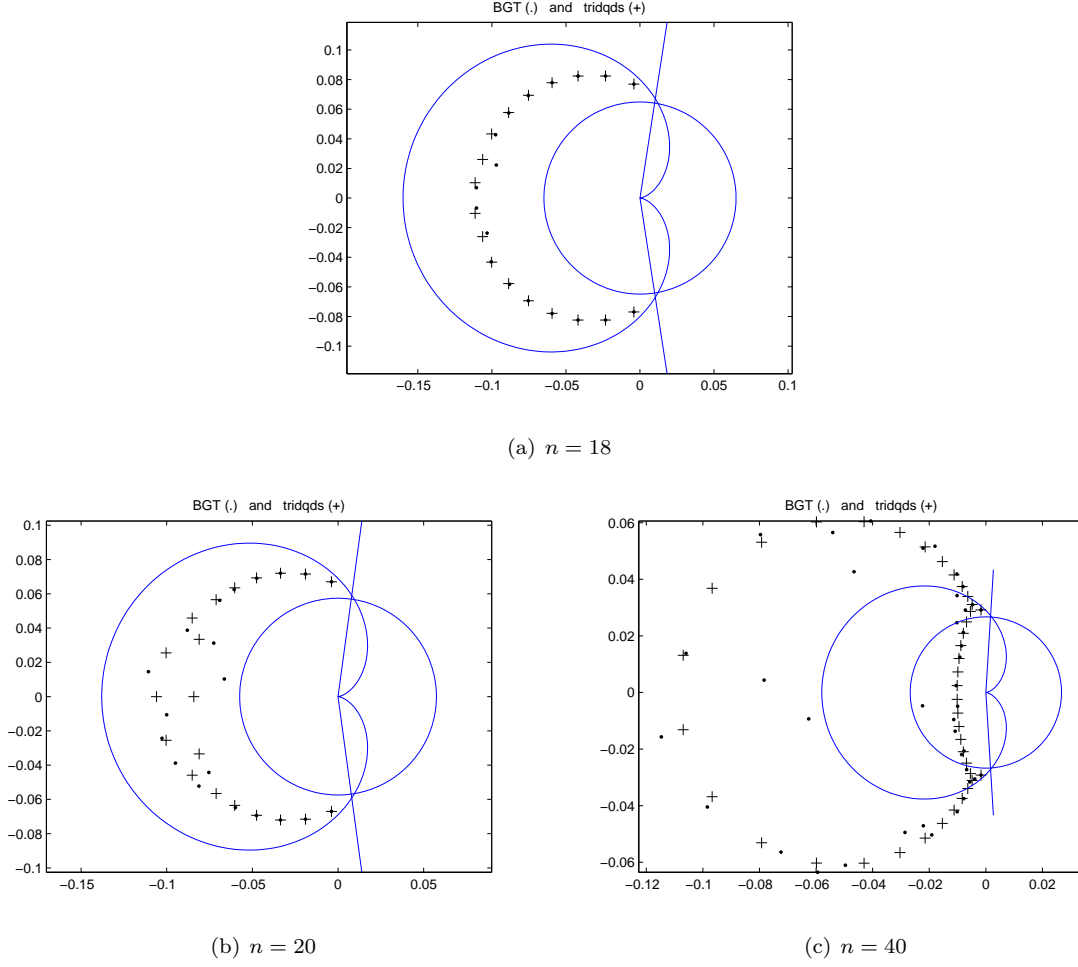
$$\alpha_1 = -\frac{b}{a}, \quad \gamma_1 = -\alpha_1, \quad \beta_1 = \frac{\alpha_1}{a+1},$$

and

$$\begin{aligned} \alpha_j &:= -b \frac{a-2}{(2j+a-2)(2j+a-4)}, \quad j = 2, \dots, n, \\ \gamma_j &:= b \frac{j+a-2}{(2j+a-2)(2j+a-3)}, \\ \beta_j &:= -b \frac{j}{(2j+a-1)(2j+a-2)}, \quad j = 2, \dots, n-1. \end{aligned}$$

Parameter b is a scaling factor and most authors take $b = 2$ and so do we. The case $a \in \mathbb{R}$ is the most investigated in literature. The eigenvalues of $B_n^{(a,b)}$, well separated complex eigenvalues, suffer from ill-conditioning that increases with n - close to a defective matrix. In Pasquini [17] it is mentioned that the ill-conditioning seems to reach its maximum when a ranges from -8.5 to -4.5 .

Our examples take $B_n^{(-4.5,2)}$ for $n = 18, 20, 40$. We show pictures for BGT and $tridqds$ to illustrate the extreme sensitivity of some of the eigenvalues. The results of complex $dqds$ are visually identical to $tridqds$, so we don't show them. In exact arithmetic the spectrum lies on an arc in the interior of the moon-shaped region. See Figure 7.1.

FIGURE 7.1. Eigenvalues of Bessel matrix $B_n^{(-4.5, 2)}$

Clement matrix

The so-called *Clement* matrices (see [3])

$$C = \text{tridiag}(\mathbf{b}, \mathbf{0}, \mathbf{c})$$

with $b_j = j$ and $c_j = b_{n-j}$, $j = 1, \dots, n-1$, have real eigenvalues

$$\begin{aligned} &\pm n-1, n-3, \dots, 1, && \text{for } n \text{ even,} \\ &\pm n-1, n-3, \dots, 0, && \text{for } n \text{ odd.} \end{aligned}$$

These matrices posed no serious difficulties. The initial zero diagonal obliges the *dqds* based methods to take care when finding an initial L, U factorization.

The *tridqds* code uses only real *dqds* transforms as it should. Our accuracy is less than *BGT* but satisfactory. The complex *dqds* and *tridqds* performed identically. The ratio of elapsed times is the striking feature.

Our numerical tests have $n = 100, 200, 400, 800$. The minimum and maximum relative errors, rel_{min} and rel_{max} , are shown in Table 7.1 and the CPU times in Table 7.2.

	<i>BGT</i>		<i>complex dqds</i>		<i>tridqds</i>	
n	rel_{min}	rel_{max}	rel_{min}	rel_{max}	rel_{min}	rel_{max}
100	0	$3 \cdot 10^{-16}$	0	$3 \cdot 10^{-14}$	0	$3 \cdot 10^{-14}$
200	0	$4 \cdot 10^{-16}$	0	$3 \cdot 10^{-13}$	0	$3 \cdot 10^{-13}$
400	0	$1 \cdot 10^{-15}$	0	$3 \cdot 10^{-12}$	0	$3 \cdot 10^{-12}$
800	0	$1 \cdot 10^{-15}$	$2 \cdot 10^{-16}$	$2 \cdot 10^{-12}$	$2 \cdot 10^{-16}$	$2 \cdot 10^{-12}$

TABLE 7.1
Relative errors for Clement matrices

n	<i>BGT</i>	<i>complex dqds</i>	<i>tridqds</i>
100	4.2	0.06	0.06
200	12.6	0.12	0.13
400	42.2	0.45	0.50
800	174.2	1.8	1.8

TABLE 7.2
CPU times in seconds for Clement matrices

Graded matrix

This matrix C was created in ΔT form with $T = \text{tridiag}(\mathbf{b}, \mathbf{a}, \mathbf{c})$,

$$a_j = b_j = c_j = 3^{-(j-1)}, \quad j = 1, \dots, n-1; \quad a_n = 3^{-(n-1)},$$


and $\Delta = \text{diag}(\boldsymbol{\delta})$, $\delta_j = (-1)^{\lfloor (j+1)/2 \rfloor}$, $j = 1, \dots, n$. The result is a balanced matrix with eigenvalues of different magnitude.

Figure 7.2 shows the approximated eigenvalues for $n = 100$. Table 7.3 reports the CPU times.

n	<i>BGT</i>	<i>complex dqds</i>	<i>tridqds</i>
50	0.31	0.02	0.02
100	0.67	0.06	0.04
200	2.12	0.14	0.06
400	...	0.45	0.35

TABLE 7.3
CPU times in seconds for the graded matrices

BGT code reported the message “Exceed maximum number of operations” for $n = 400$. The performance of all methods for the flipped matrix is practically the same.

FIGURE 7.2. *Eigenvalues of the graded matrix with $n = 100$*

Matrix with clusters

Matrix Test 5 in [1],

$$C = D^{-1} \text{tridiag}(\mathbf{1}, \boldsymbol{\alpha}, \mathbf{1}), \quad D = \text{diag}(\boldsymbol{\beta}), \quad \boldsymbol{\alpha}, \boldsymbol{\beta} \in \mathbb{R}^n$$

$$\alpha_k = 10^{5(-1)^k} \cdot (-1)^{\lfloor k/4 \rfloor}, \quad \beta_k = (-1)^{\lfloor k/3 \rfloor}, \quad k = 1, \dots, n,$$

seems to be a challenging test matrix. It was designed to have large, tight clusters of eigenvalues around 10^{-5} , -10^5 and 10^5 . Half the spectrum is around 10^{-5} and the rest is divided unevenly between -10^5 and 10^5 . The diagonal alternates between entries of absolute value 10^5 and 10^{-5} and so, for *dqds* codes, there is a lot of rearranging to do. When $n \geq 100$ it is not clear what is meant by accuracy.

All three codes obtain the correct number of eigenvalues in each cluster and the diameters of the clusters are all about 10^{-5} . The striking feature is the time taken. See Table 7.4.

n	<i>BGT</i>	<i>complex dqds</i>	<i>tridqds</i>
50	1.2	0.03	0.01
100	4.5	0.05	0.03
200	20.1	0.14	0.08
400	85.0	0.61	0.13

TABLE 7.4
CPU times in seconds for Test 5 matrix

Matrix Test 4

For matrix Test 4 in [1],

$$C = D^{-1} \text{tridiag}(\mathbf{1}, \boldsymbol{\alpha}, \mathbf{1}), \quad D = \text{diag}(\boldsymbol{\beta}), \quad \boldsymbol{\alpha}, \boldsymbol{\beta} \in \mathbb{R}^n$$

$$\alpha_k = (-1)^k, \quad \beta_k = 20 \cdot (-1)^{\lfloor k/5 \rfloor}, \quad k = 1, \dots, n,$$

the performance of the three codes is shown in Table 7.5. Figure 7.3 shows the eigenvalues of this matrix for $n = 50$.

n	<i>BGT</i>	<i>complex dqds</i>	<i>tridqds</i>
50	1.0	0.05	0.03
100	4.3	0.08	0.03
200	17.8	0.27	0.09
400	78.6	1.0	0.44
800	342.6	5.5	2.4

TABLE 7.5
CPU times in seconds for Test 4 matrix




FIGURE 7.3. *Eigenvalues of Test 4 matrix with $n = 50$*

Liu matrix

Z. A. Liu [13] devised an algorithm to obtain one-point spectrum unreduced tridiagonal matrices of arbitrary dimension $n \times n$. These matrices have only one eigenvalue, zero with multiplicity n , and the Jordan form consists of one Jordan block. Our code *tridqds* computes this eigenvalue exactly (and also the generalized eigenvectors) using the following method which is part of the prologue.

The best place to start looking for eigenvalues of a tridiagonal matrix $C = \text{tridiag}(\mathbf{a}, \mathbf{b}, \mathbf{c})$ is at the arithmetic mean which we know ($\mu = \text{trace}(C)/n$). Before converting to J -form and factoring, we check whether μ is an eigenvalue by using the 3-term recurrence to solve

$$(\mu I - C)\mathbf{x} = \mathbf{e}_n p_n(\mu) / \prod_{i=1}^{n-1} c_i.$$

Here

$$\begin{aligned} x_1 &= 1, & x_2 &= (\mu - a_2)/c_1, \\ x_{j+1} &= \frac{1}{c_j} [(\mu - a_j)x_j - b_{j-1}x_{j-1}], & j &= 2, \dots, n-1, \end{aligned}$$

and

$$v := (\mu - a_n)x_n - b_{n-1}x_{n-1} \left(= p_n(\mu) / \prod_{i=1}^{n-1} c_i \right).$$

If, by chance, v vanishes, or is negligible compared to $\|\mathbf{x}\|$, then μ is an eigenvalue (to working accuracy) and \mathbf{x} is an eigenvector. To check its multiplicity we differentiate with respect to μ and solve

$$(\mu I - C)\mathbf{y} = \mathbf{x}$$

with $y_1 = 0$, $y_2 = 1 = x'_2 (= x_1)$. If

$$v' = p'_n(\mu) / \prod_{i=1}^{n-1} c_i := (\mu - a_n)y_n - b_{n-1}y_{n-1} + x_n$$

vanishes, or is negligible w.r.t. $\|\mathbf{y}\|$, then we continue the same way until the system is inconsistent or there are n generalized eigenvectors.

Usually $v \neq 0$ and the calculation appears to have been a waste. This is not quite correct. In exact arithmetic, triangular factorization of $\mu I - C$ or $\mu I - J$, where $J = \Delta C \Delta^{-1}$, will break down if, and only if, x_j vanishes for $1 < j < n$. So our code examines $\min_j |x_j|$ and if it is too small w.r.t. its neighbors and w.r.t. $\|\mathbf{x}\|$ then we do not choose μ as our initial shift. Otherwise we do obtain initial L and U from $J - \mu I = LU$.

8. Conclusions and future work. We conclude that, working together, a single *dqds* transform with real shifts and our *tridqds* transform with complex conjugate pairs of shifts constitute the right tool for computing the eigenvalues of real tridiagonal matrices.

However there is far more work to be done for the following reasons. In a previous paper we discovered that, surprisingly often, eigenvalues are determined to, not high, but adequate relative accuracy; tiny relative changes η in the parameters that define the matrix produce relative changes in the eigenvalue of the order of $10^3\eta$ or $10^4\eta$. This is good news. We cannot tell in advance when this occurs. In our opinion a relative condition number should be returned with each eigenvalue. This requires an approximation to the row and column eigenvectors, whether or not the user needs them.

We envision software that computes an initial approximation to each eigenvalue and then invokes a generalized Rayleigh quotient iteration to both compute eigenvectors and obtain a refined

eigenvalue approximation, along with the smallest residual norms that could be achieved. Then the relative condition number can be formed. See [9].

Another practical feature is to scan the initial matrix to extract Gersgorin disks and a tight box in the complex plane that contains the spectrum. Matrices from industrial sources frequently permit “localization” of the eigenvectors belonging to certain parts of the spectrum. One consequence is that the relevant eigenvectors, and eigenvalues, may be obtained from small submatrices.

It is also important to scale and normalize the initial matrix and make use of the splitting that occurs with big matrices. Currently our shift strategy is quite straightforward and it is both difficult and worthwhile to improve it. There are plenty of challenges to be met before software for this real tridiagonal problem can be installed in packages such as LAPACK, not to mention parallel computation and scalapack.

Appendix A. Implementation details of minor step i . In this appendix we show how the calculations involved in each minor step of *tridqds* can be organized.

For each minor step i , $i = 2, \dots, n-3$, consider F and G as in (4.3). Denote the 2×1 bulge in F , indicated with plus signs, by $[x_l \ y_l]^T$. And denote the entries $G_{i,i}$, $G_{i+1,i}$ and $G_{i+2,i}$, indicated with $+$, $+$, $+$, by $[x_r \ y_r \ z_r]^T$. Subscripts l and r derive from “left” and “right”, respectively. This way we have

$$F = \begin{bmatrix} \ddots & & & & & \\ \ddots & 1 & & & & \\ & \hat{l}_{i-1} & u_i & 1 & & \\ & x_l & & u_{i+1} & 1 & \\ & y_l & & & u_{i+2} & \ddots \\ & & & & & \ddots \end{bmatrix}, \quad G = \begin{bmatrix} \ddots & & \ddots & & & \\ & \hat{u}_{i-1} & 1 & & & \\ & & x_r & & & \\ & & y_r & 1 & & \\ & & z_r & l_{i+1} & 1 & \\ & & & & \ddots & \ddots \end{bmatrix}$$

and the minor step i can be accomplished using only these auxiliary variables.

Minor step i

- a) • Matrices Z_i^{-1} and Z_i

$$Z_i^{-1} = \begin{bmatrix} \ddots & & & & & \\ & 1 & & & & \\ & & \frac{1}{u_i} & -\frac{1}{u_i} & & \\ & & 0 & 1 & & \\ & & & & 1 & \\ & & & & & \ddots \end{bmatrix}, \quad Z_i = \begin{bmatrix} \ddots & & & & & \\ & 1 & & & & \\ & & u_i & 1 & & \\ & & 0 & 1 & & \\ & & & & 1 & \\ & & & & & \ddots \end{bmatrix}$$

- The effect of Z_i^{-1} and the effect of Z_i

$$FZ_i^{-1} = \begin{bmatrix} \ddots & & & & & \\ \ddots & 1 & & & & \\ & \hat{l}_{i-1} & 1 & 0 & & \\ & x_l & & u_{i+1} & 1 & \\ & y_l & & & u_{i+2} & \ddots \\ & & & & & \ddots \end{bmatrix}, Z_i G = \begin{bmatrix} \ddots & & & & & \\ & \ddots & & & & \\ & \hat{u}_{i-1} & 1 & & & \\ & & x_r & 1 & & \\ & & y_r & 1 & & \\ & & z_r & l_{i+1} & 1 & \\ & & & & \ddots & \ddots \end{bmatrix}$$

where

$$x_r \leftarrow x_r * u_i + y_r$$

- b) • Matrices \mathcal{L}_i^{-1} and \mathcal{L}_i

$$\mathcal{L}_i^{-1} = \begin{bmatrix} \ddots & & & & & \\ & \ddots & & & & \\ & & 1 & & & \\ & & x_l & 1 & & \\ & & y_l & & 1 & \\ & & & & & \ddots \end{bmatrix}, \quad \mathcal{L}_i = \begin{bmatrix} \ddots & & & & & \\ & \ddots & & & & \\ & & 1 & & & \\ & & -x_l & 1 & & \\ & & -y_l & & 1 & \\ & & & & & \ddots \end{bmatrix}$$

where

$$\begin{aligned} x_l &\leftarrow -x_l/\hat{l}_{i-1} \\ y_l &\leftarrow -y_l/\hat{l}_{i-1} \end{aligned}$$

- The effect of \mathcal{L}_i^{-1}

$$\mathcal{L}_i^{-1} F = \begin{bmatrix} \ddots & & & & & \\ \ddots & 1 & & & & \\ & \hat{l}_{i-1} & 1 & & & \\ & & x_l & u_{i+1} & 1 & \\ & & y_l & & u_{i+2} & \ddots \\ & & & & & \ddots \end{bmatrix}$$

- The effect of \mathcal{L}_i

$$G\mathcal{L}_i = \begin{bmatrix} \ddots & & & & & \\ & \ddots & & & & \\ & \hat{u}_i & 1 & & & \\ & x_r & 1 & & & \\ & y_r & l_{i+1} & 1 & & \\ & z_r & & l_{i+2} & 1 & \\ & & & & \ddots & \ddots \end{bmatrix}$$

where

$$\begin{aligned}\hat{u}_i &\leftarrow x_r - x_l \\ x_r &\leftarrow y_r - x_l \\ y_r &\leftarrow z_r - y_l - x_l * l_{i+1} \\ z_r &\leftarrow -y_l * l_{i+2}\end{aligned}$$

- c) • Matrices Y_i^{-1} and Y_i

$$Y_i^{-1} = \begin{bmatrix} \ddots & & & & \\ & 1 & & & \\ & x_r & 1 & & \\ & y_r & & 1 & \\ & z_r & & & 1 \\ & & & & & \ddots \end{bmatrix}, \quad Y_i = \begin{bmatrix} \ddots & & & & \\ & 1 & & & \\ & -x_r & 1 & & \\ & -y_r & & 1 & \\ & -z_r & & & 1 \\ & & & & & \ddots \end{bmatrix}$$

where

$$\begin{aligned}x_r &\leftarrow x_r / \hat{u}_i \\ y_r &\leftarrow y_r / \hat{u}_i \\ z_r &\leftarrow z_r / \hat{u}_i\end{aligned}$$

- The effect of Y_i^{-1}

$$FY_i^{-1} = \begin{bmatrix} \ddots & & & & \\ & \ddots & & & \\ & & 1 & & \\ & & \hat{l}_i & u_{i+1} & 1 \\ & & x_l & & u_{i+2} & 1 \\ & & & & & & 1 \\ & & & & & & & u_{i+3} & \ddots \\ & & & & & & & & \ddots \end{bmatrix}$$

where

$$\begin{aligned}\hat{l}_i &\leftarrow x_l + y_r + x_r * u_{i+1} \\ x_l &\leftarrow y_l + z_r + y_r * u_{i+2} \\ y_l &\leftarrow z_r * u_{i+3}\end{aligned}$$

- The effect of Y_i

$$Y_i G = \begin{bmatrix} \ddots & & & & \\ & \ddots & & & \\ & & \hat{u}_i & 1 & \\ & & & x_r & \\ & & & y_r & 1 \\ & & & z_r & l_{i+2} & 1 \\ & & & & & & \ddots & \ddots \end{bmatrix}$$

where

$$\begin{aligned}x_r &\leftarrow 1 - x_r \\ y_r &\leftarrow l_{i+1} - y_r \\ z_r &\leftarrow -z_r\end{aligned}$$

Appendix B. *tridqds* algorithm.

tridqds($\sigma, \bar{\sigma}$) :

```

% step 1
 $x_r = 1$ ;  $y_r = l_1$ ;  $z_r = 0$ 
% the effect of  $Z_1$ 
 $x_r = x_r * u_1 + y_r$ 
% the matrix  $\mathcal{L}_1^{-1}$ 
 $x_l = (u_1 + l_1)^2 + u_2 l_1 - 2(\Re \sigma)(u_1 + l_1) + |\sigma|^2$ 
 $y_l = -u_2 l_1 u_3 l_2 / x_l$ 
 $x_l = -u_2 l_1 (u_1 + l_1 + u_2 + l_2 - 2(\Re \sigma)) / x_l$ 
% the effect of  $\mathcal{L}_1$ 
 $\hat{u}_1 = x_r - x_l$ ;
 $x_r = y_r - x_l$ ;  $y_r = z_r - y_l - x_l * l_2$ ;  $z_r = -y_l * l_3$ 
% the matrix  $Y_1^{-1}$ 
 $x_r = x_r / \hat{u}_1$ ;  $y_r = y_r / \hat{u}_1$ ;  $z_r = z_r / \hat{u}_1$ 
% the effect of  $Y_1^{-1}$ 
 $\hat{l}_1 = x_l + y_r + x_r * u_2$ 
 $x_l = y_l + z_r + y_r * u_3$ ;  $y_l = z_r * u_4$ 
% the effect of  $Y_1$ 
 $x_r = 1 - x_r$ ;  $y_r = l_2 - y_r$ ;  $z_r = -z_r$ 

for  $i = 2, \dots, n - 3$ 
% the effect of  $Z_i$ 
 $x_r = x_r * u_i + y_r$ 
% the matrix  $\mathcal{L}_i^{-1}$ 
 $x_l = -x_l / \hat{l}_{i-1}$ ;  $y_l = -y_l / \hat{l}_{i-1}$ ;
% the effect of  $\mathcal{L}_i$ 
 $\hat{u}_i = x_r - x_l$ ;
 $x_r = y_r - x_l$ ;  $y_r = z_r - y_l - x_l * l_{i+1}$ ;  $z_r = -y_l * l_{i+2}$ 
% the matrix  $Y_i^{-1}$ 
 $x_r = x_r / \hat{u}_i$ ;  $y_r = y_r / \hat{u}_i$ ;  $z_r = z_r / \hat{u}_i$ 
% the effect of  $Y_i^{-1}$ 
 $\hat{l}_i = x_l + y_r + x_r * u_{i+1}$ 
 $x_l = y_l + z_r + y_r * u_{i+2}$ ;  $y_l = z_r * u_{i+3}$ 
% the effect of  $Y_i$ 
 $x_r = 1 - x_r$ ;  $y_r = l_{i+1} - y_r$ ;  $z_r = -z_r$ 
end for

```

```

% step n-2
% the effect of  $Z_{n-2}$ 
 $x_r = x_r * u_{n-2} + y_r$ 
% the matrix  $\mathcal{L}_{n-2}^{-1}$ 
 $x_l = -x_l / \hat{l}_{n-3}; \quad y_l = -y_l / \hat{l}_{n-3};$ 
% the effect of  $\mathcal{L}_{n-2}$ 
 $\hat{u}_{n-2} = x_r - x_l;$ 
 $x_r = y_r - x_l; \quad y_r = z_r - y_l - x_l * l_{n-1}$ 
% the matrix  $Y_{n-2}^{-1}$ 
 $x_r = x_r / \hat{u}_{n-2}; \quad y_r = y_r / \hat{u}_{n-2}$ 
% the effect of  $Y_{n-2}^{-1}$ 
 $\hat{l}_{n-2} = x_l + y_r + x_r * u_{n-1}$ 
 $x_l = y_l + y_r * u_n$ 
% the effect of  $Y_{n-2}$ 
 $x_r = 1 - x_r; \quad y_r = l_{n-1} - y_r$ 

```

```

% step n-1
% the effect of  $Z_{n-1}$ 
 $x_r = x_r * u_{n-1} + y_r$ 
% the matrix  $\mathcal{L}_{n-1}^{-1}$ 
 $x_l = -x_l / \hat{l}_{n-2}$ 
% the effect of  $\mathcal{L}_{n-1}$ 
 $\hat{u}_{n-1} = x_r - x_l;$ 
 $x_r = y_r - x_l$ 
% the matrix  $Y_{n-1}^{-1}$ 
 $x_r = x_r / \hat{u}_{n-1}$ 
% the effect of  $Y_{n-1}^{-1}$ 
 $\hat{l}_{n-1} = x_l + x_r * u_n$ 
% the effect of  $Y_{n-1}$ 
 $x_r = 1 - x_r$ 

```

```

% step n
% the effect of  $Z_n$ 
 $x_r = x_r * u_n$ 
% the matrix  $\mathcal{L}_n^{-1} = I$ 
% the effect of  $\mathcal{L}_n$ 
 $\hat{u}_n = x_r;$ 
% the matrix  $Y_n^{-1} = I$ 

```

REFERENCES

- [1] D. A. Bini, L. Gemignani and F. Tisseur. *The Ehrlich-Aberth method for the nonsymmetric tridiagonal eigenvalue problem*. SIAM J. Matrix Analysis and its Applications, 27(1):153-175, 2005.
- [2] A. Bunse-Gerstner. *An analysis of the HR algorithm for computing the eigenvalues of a matrix*. Linear Algebra and Its Applications, 35:155-173, 1981.
- [3] P. A. Clement. *A class of triple-diagonal matrices for test purposes*. SIAM Review, 1 (vol.1), January, 1959.
- [4] Jane K. Cullum. *A QL procedure for computing the eigenvalues of complex symmetric tridiagonal matrices*. SIAM J. Matrix Analysis and its Applications, 17(1):83-109, 1996.
- [5] D. Day. *Semi-duality in the two-sided lanczos algorithm*. Ph.D thesis, University of California, Berkeley, 1993.
- [6] I. S. Dhillon and B. N. Parlett. *Multiple representations to compute orthogonal eigenvectors of symmetric tridiagonal matrices*. Linear Algebra and its Applications, 387:1-28, 2004.
- [7] K. Fernando and B. Parlett. *Accurate singular values and differential qd algorithms*. Numerische Mathematik, 67:191-229, 1994.
- [8] C. Ferreira. *The unsymmetric tridiagonal eigenvalue problem*. Ph.D Thesis, University of Minho, 2007. <http://hdl.handle.net/1822/6761>
- [9] C. Ferreira and B. Parlett. *Sensitivity of Eigenvalues of an unsymmetric tridiagonal matrix*. Submitted.
- [10] C. Ferreira and B. Parlett. *Convergence of LR algorithm for a one-point spectrum tridiagonal matrix*. Numerische Mathematik, 113(3):417-431, 2009.
- [11] C. Ferreira and B. Parlett. *Sensitivity of Eigenvalues of an unsymmetric tridiagonal matrix*. Submitted.
- [12] J. G. F. Francis. *The QR transformation - a unitary analogue to the LR transformation, Parts I and II*. Computer Journal, 4:265-272 and 332-245, 1961/62.
- [13] Z. A. Liu. *On the extended HR algorithm*. Technical Report PAM-564, Center for Pure and Applied Mathematics, University of California, Berkeley, CA, USA, 1992.
- [14] B. N. Parlett. *Reduction to tridiagonal form and minimal realizations*, SIAM Journal on Matrix Analysis, 13:567-593, 1992.
- [15] B. N. Parlett. *The new qd algorithms*. Acta Numerica, 459-491, 1995.
- [16] B. N. Parlett and Osni A. Marques. *An implementation of the dqds algorithm*. Linear Algebra and Its Applications, 309:217-259, 2000.
- [17] L. Pasquini. *Accurate computation of the zeros of the generalized Bessel polynomials*. Numerische Mathematic, 86:507-538, 2000.
- [18] H. Rutishauser. *Solution of eigenvalue problems with the LR-transformation*. National Bureau of Standards Applied Mathematics series 49:47-81, 1958.
- [19] H. Rutishauser and H.R. Schwarz. *The LR transformation method for symmetric matrices*. Numerische Mathematik, 5:273-289, 1963.
- [20] J. Slemons. *The Result of Two Steps of the LR Algorithm is Diagonally Similar to the Result of One Step of the HR Algorithm*. SIAM J. Matrix Analysis and its Applications, 31(1):68-74, 2009.
- [21] D. S. Watkins. *QR-like algorithms - An overview of convergence theory and practice*. Lectures in Applied Mathematics, 32:879-893, 1996.
- [22] D. S. Watkins and L. Elsner. *Convergence of algorithms of decomposition type for the eigenvalue problem*. Linear Algebra and Its Applications, 143:19-47, 1991.
- [23] Z. Wu. *The Triple dqds Algorithm for Complex Eigenvalues*. Ph.D thesis, University of California, Berkeley, 1996.
- [24] H. Xu. *The relation between the QR and LR algorithms*. SIAM J. Matrix Analysis and its Applications, 19(2):551-555, 1998.
- [25] Yao Yang. *Error Analysis of the qds and dqds Algorithms*. Ph.D thesis, University of California, Berkeley, 1994.